

PCTWORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

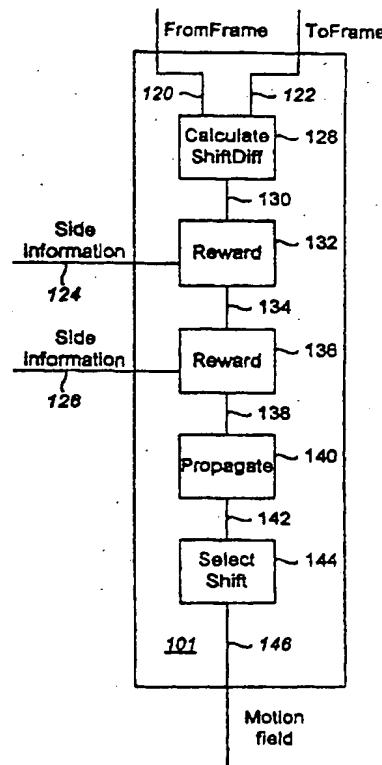
(51) International Patent Classification ⁶ :	A1	(11) International Publication Number:	WO 95/26539
G06T 7/20		(43) International Publication Date:	5 October 1995 (05.10.95)

(21) International Application Number:	PCT/EP95/01011			
(22) International Filing Date:	17 March 1995 (17.03.95)			
(30) Priority Data:	941126	25 March 1994 (25.03.94)	NO	
(71) Applicant (for all designated States except US):	IDT INTERNATIONAL DIGITAL TECHNOLOGIES DEUTSCHLAND GMBH [DE/DE]; Fraunhoferstrasse 9, D-86737 Ismaning (DE).			
(72) Inventors; and				
(75) Inventors/Applicants (for US only):	MARTENS, Harald, Aa-gaard [NO/DE]; Freischützstrasse 84, D-81927 München (DE). REBERG, Jan, Otto [NO/DE]; Fritz-Meyer-Weg 55, C 4/8, D-81925 München (DE).			
(74) Agent:	BETTEN & RESCH; Reichenbachstrasse 19, D-80469 München (DE).			

(54) Title: METHOD AND APPARATUS FOR ESTIMATING MOTION

(57) Abstract

The present invention provides a method and apparatus for estimating motion between at least two frames, each consisting of a plurality of samples of an input signal, the method comprising the steps of: (1) selecting for each sample in the first frame a plurality of candidate shifts, (2) computing for each candidate shift a measurement of fit between said sample in said first frame and the corresponding sample at the shifted position in said second frame, (3) selecting for each sample in said first frame the shift with the best associated fit, wherein the collection of selected shifts represents the motion from said first frame to said second frame, and, prior to step (3), (4) adjusting for each sample in said first frame the measurement of fit according to an externally available hypothesis indicating a priori knowledge about how probable it is that certain shifts represent the correct motion.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgyzstan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Latvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam
GA	Gabon				

DESCRIPTION

METHOD AND APPARATUS FOR ESTIMATING MOTION

Field of the Invention

The present invention relates to a method and apparatus for estimating or analyzing movement of data patterns between similar frames. The invention can be used for estimating motion between video images, for recognizing voice, and for a number of other applications in connection with physical entities moving in time, space, frequency, intensity, color, shape, etc.

In image analysis and compression, this is often called motion estimation, motion vector selection, disparity analysis, or finding optic flow or velocity patterns. In signal processing it is called dynamic time warping.

Background of the Invention

A general description of this field is given by Ajit Singh ("Optic Flow Computation", IEEE Computer Society Press, 1991).

EP-A-348 207 describes an image motion vector detecting apparatus. US-A-5 276 513 describes an implementation architecture for performing hierarchical motion analysis of video images in real time.

EP-A-538 667 describes an adaptive motion compensation using a plurality of motion compensators.

PCT/US94/10190 describes a method and apparatus for data analysis.

All references mentioned before are incorporated herein by reference.

There already exist methods according to the above references. But these methods have some problems:

Too simplistic models

In video compression systems, like MPEG and MPEG II, one common assumption is that one frame usually can be reconstructed by dividing the frame into blocks, often consisting of 16*16 pixels, and then for each block locating and transmitting position of the most similar block in one or two already transmitted frames. This method is reasonably simple to implement, but it does not necessarily provide optimal compression ratios.

Noise robustness

When the result is based on only one pair of data sets, then noise in the input data sets will result in noise in the output.

Precision

When the result is based on only one pair of data sets, this limits the precision of the output. E.g., for matching digital images, it is impossible without any preassumptions to find the match between two pixels with better than one pixel accuracy in position.

Ambiguity

When matching one feature in one frame against another frame, sometimes there are more than one possible match. This creates problems for many methods.

Accordingly, it is an object of the present invention to provide a method and apparatus for estimating movement of data patterns between similar frames which solve the above mentioned problems.

It is also an object of the present invention to provide a method and apparatus for estimating movement of data patterns between similar frames which provide efficient and easy detection of the motion of moving data patterns.

Summary of the Invention

This invention is mainly concerned with the processing of at least two frames, called "FromFrame" and "ToFrame", finding for each sample in the first frame (FromFrame) where the corresponding sample is located in the second frame (ToFrame), so that the "motion field" could efficiently be detected.

A first embodiment of the present invention provides a method and apparatus for estimating motion between at least two frames, each consisting of a plurality of samples of an input signal, the method comprising the steps of: (1) selecting for each sample in the first frame a plurality of candidate shifts, (2) computing for each candidate shift a measurement of fit between said sample in said first frame and the corresponding sample at the shifted position in said second frame, (3) selecting for each sample in said first frame the shift with the best associated fit, wherein the collection of selected shifts represents the motion from said first frame to said second frame, and, prior to step (3), (4) adjusting for each sample in said first frame the measurement of fit according to an externally available hypothesis indicating a priori knowledge about how probable it is that certain shifts represent the correct motion.

A further embodiment of the present invention provides a method and apparatus for estimating motion between at least two frames each consisting of a plurality of samples of an input signal, the method comprising the steps of: (1) selecting for each sample in the first frame a plurality of candidate shifts, (2) computing for each candidate shift a measurement of fit between said sample in said first frame and

the corresponding sample at the shifted position in said second frame, (3) selecting for each sample in said first frame the shift with the best associated fit, wherein the collection of selected shifts represents the motion from said first frame to said second frame, and, prior to step (3), (4) adjusting for each sample in said first frame the measurement of fit for each shift according to measurements of fit for the same shifts for adjacent samples in said first frame.

The present invention provides an improvement of existing technology by exploiting more data and a priori knowledge in the analysis. It describes a method and apparatus for assembling together information from several sources, e.g.

- Preceeding or otherwise similar frames:

Instead of using only two frames as basis for the estimation, this invention allows information from several frames to be exploited at the same time.

- Other sensors or methods:

When collecting raw data, one often has the possibility of using several methods. E.g., when preparing terrain maps, one could combine aerial photographs with traditional triangulation or Global Positioning System (GPS). This invention makes it possible to combine several such different data sets into one common framework.

- Regularity/smoothness:

One often has some knowledge of how smooth the solution should be. E.g., when preparing terrain maps, one usually knows if the area consists of mountains or fields. This invention gives a possibility of exploiting such knowledge.

Brief Description of the Drawings

Fig. 1a gives an overview of the method according to prior art;

Fig. 1b gives an overview of the method according to the invention using a motion estimator according to the invention;

Fig. 2 gives a numeric example;

Fig. 3 shows a pyramidal framework;

Fig. 4 shows a pipelined architecture, suitable for the pyramidal framework in Fig. 3;

Fig. 5 shows in more detail the inner workings of Reward; and

Fig. 6 shows in more detail the computation of slacks shown in Fig. 2.

First of all, some terms used in this description will be explained.

Definitions

One set of samples of an input signal from one domain, which could be a temporal, spatial, color or shape domain, will here be called a frame. For an example from sound applications, a frame could consist of air pressure measurements, when air pressure varying as a function of time is picked up with a microphone and recorded on some storage medium, during utterance of one word. It would be practical to address the samples using one number, corresponding in this example to time, and so this structure will be called a one-dimensional frame, or 1D frame for short.

For an example from imaging applications, a frame could consist of light intensity measurements. Such samples are usually called picture elements, or also just "pixels". Here it would be practical to address the samples using two numbers, namely one indicating which scan line a pixel belongs to (which will be called its vertical position) and one indicating its position along this scan line (which will be called its horizontal position), and so this structure will be called a two-dimensional frame, or 2D frame for short.

For an example from higher-dimensional imaging applications, a frame could consist of tomography samples from a MRI, PET or similar 3-dimensional volume scanner. The samples correspond to some radiation attenuation or emission caused by the measured object. Each sample is addressed using a combination of up to three spatial coordinates and possibly also a temporal coordinate.

The term "motion" will be used when similar patterns have different "positions" in the frames. This will not necessarily be the same as motion in the traditional sense; e.g., for sound applications, this "motion" can be interpreted as shift in time or varying delay.

In order to more readily understand the invention a brief description of prior art is given with reference to Fig. 1a.

Prior art can be summarized with reference to Fig. 1a showing a motion estimator 100. A FromFrame 102 and a ToFrame 104 are input to a module 106 for calculating a Shift Difference 108 (to be defined below). A further module 110 selects for each element in the FromFrame a "best" shift, and together this collection of shifts constitutes a motion field 112.

Block based techniques, like used in MPEG image sequence compression standards, divide FromFrame 102 into blocks and for each block evaluate different shifts into ToFrame 104, and the module for selecting shift 110 usually selects shifts minimizing sum of absolute value of displaced frame difference.

Gradient based methods like these presented by B.K.P. Horn/B.Schunck ("Determining Optical Flow", Artificial Intelligence 17, 1981), which is incorporated herein by reference, first calculate ShiftDiff, together with gradients, according to some motion hypothesis, and then use this information for selecting a better motion hypothesis. This is iterated, producing better motion fields. A variation of this technique with fewer or no iterations is described in EP-A-596 409.

Phase correlation based methods perform an analysis in the frequency domain in order to select only the few most probable candidate shifts, before ShiftDiff is computed and the "best" shifts are selected. WO-A-87/05769 describes a TV picture motion measurement system based on this principle.

Feature based methods first extract features from the frames, then calculate ShiftDiff and select shifts for these features. The output motion for the features can then be interpolated or otherwise converted to apply for the whole frame if wanted.

Detailed Description of the Preferred Embodiments

A general overview of the invention gives Fig. 1b showing a motion estimator 101, in short EstMov 101, according to the invention. According to prior art EstMov 101 consists of a module 128 for calculating the Shift Difference 130 of two input frames 120, 122 and a module 144 for selecting shifts, and also has the novel feature of being able to exploit side information 124 and 126 through one or more reward modules 132, 136 and the novel feature of a module 140 for propagating information in the ShiftDifference.

The side information consists of hypotheses for a motion field, and optionally also strengths and ranges for these hypotheses, which is further explained later.

The invention will be exemplified by three different preferred embodiments. The first will be for 1D frames, which is less difficult to explain.

First Preferred Embodiment

The method will be explained with reference to a numerical example given in Fig. 2. The motion is to be computed from a FromFrame 210, consisting of 6 elements, to a ToFrame 212, consisting of 7 elements. By visual inspection, it seems that the elements in FromFrame have moved one step to the right. For this example, it is also given that based on previous results, e.g. from analysis of previous frames if this were two frames from a sequence, one has found that there is probably no motion in the rightmost part of the frame.

It should be understood that FromFrame or ToFrame could be measured and digitized values from some sensor, or synthesized values from some model. FromFrame or ToFrame can also have been preprocessed. One useful way of preprocessing is

filtering. Low pass filtering can be utilized to disregard changes in the smallest details, which can often be modelled more efficiently with intensity modelling than with motion modelling. High pass filtering can be utilized to account for general intensity changes. This can also be done in one operation by band pass filtering. FromFrame or ToFrame may also have been subsampled.

The basic data structure to be operated on is a so-called "shift difference matrix", or for short ShiftDiff. For 1D frames, this matrix has 2 dimensions. One dimension corresponds to the number of elements in FromFrame, and the other dimension corresponds to the number of different positions in ToFrame an element of FromFrame may be moved to.

Each such position selected to be represented in ShiftDiff and thereby considered in the further analysis is called a candidate shift. One way of selecting candidate shifts for an element in FromFrame is to select all shifts up to a specified limit, which may be called SearchRange, away from its position in FromFrame.

Each element in ShiftDiff corresponds to the lack of fit when moving an element in FromFrame to the position in ToFrame. More formally, the method for calculating ShiftDiff can be defined as:

ShiftDiff:

For each Element in FromFrame:

For each considered Shift into some position in ToFrame:

ShiftDiff[Shift,Element] =

$abs(FromFrame[Element] - ToFrame[Element+Shift])$

In this preferred embodiment, absolute value of difference is used as the measure of lack of fit. It should be understood that other measures could be used. One possibility is to use square of difference.

In this preferred embodiment, the computation is performed with regard to each of the elements of FromFrame and ToFrame. It should be understood that the computation

could be performed with regard to blocks of elements. For an example on how this invention can be applied with blocks, the method for calculating ShiftDiff is given for blocks:

ShiftDiff:

For each Block of elements in FromFrame:

For each considered Shift into some position in ToFrame:

ShiftDiff[Shift,Block] =

sum(For each Element in Block:

abs(FromFrame[Element] - ToFrame[Element+Shift]))

The computed motion field will then have the resolution given by the block size.

EP-A-294 961 uses ShiftDiff according to the above definition.

The candidate shifts are not restricted to one pixel position step as indicated above. In particular, they can be smaller than one pixel position step. In this case, interpolation in ToFrame is implied. E.g, if motion is to be computed to half pixel resolution, then the method for calculating ShiftDiff should be changed into:

ShiftDiff:

For each Element in FromFrame:

For each considered Shift in half-pixel steps into some position in ToFrame:

If the current Shift refers to an integer pixel position:

ToFrameSample = ToFrame[Element + Shift/2]

else

ToFrameSample = (ToFrame[Element+(Shift-1)/2] +

ToFrame[Element + (Shift+1)/2]) / 2

ShiftDiff[Shift,Element] =

abs(FromFrame[Element] - ToFrameSample)

ShiftDiff can be based on more than two frames. E.g, when each frame corresponds to a point in time, motion with little acceleration is wanted, and several input frames are

available, it is possible to use one frame as FromFrame, the following frame as ToFrame, but then also to use the frame previous to FromFrame, in short called PrevFrame, as basis for computing ShiftDiff. It can then be assumed that if an element can correctly be described by selecting one shift representing motion from FromFrame to ToFrame, it should also be correctly described by the opposite shift from FromFrame to PrevFrame. The direction is opposite because the motion is assumed to be proportional to time and the time from FromFrame to PrevFrame is the negative of the time from FromFrame to ToFrame.

Here is how to compute ShiftDiff based on FromFrame, ToFrame and PrevFrame:

ShiftDiff:

For each Element in FromFrame:

For each candidate Shift into some position in ToFrame:

ShiftDiff[Shift,Element] =

$$\text{abs}(\text{FromFrame}[Element] - \text{FollowFrame}[Element + Shift]) +$$

$$\text{abs}(\text{FromFrame}[Element] - \text{PrevFrame}[Element - Shift])$$

Often, one has reason to assume that a rough estimate of the motion is known, and only the correction to the estimate is wanted. If at the same time the search range is large, then the above method of calculating ShiftDiff becomes inefficient, in that large parts of ShiftDiff will correspond to shifts that are not of interest. It then becomes more efficient to only select a few candidate shifts at or adjacent to the given estimate and only work on these. At the same time, it will usually be preferable to have a regular structure of ShiftDiff, in order to simplify processing. One compromise is to only have a fixed number of candidate shifts, calculated according to an offset Off. So, for each element in FromFrame, one does not necessarily compare this to elements in the same or and adjacent position in ToFrame, but rather to elements at or adjacent to a given offset position in ToFrame.

In Fig. 2, an Offset 214 is zero for all elements, for reasons of simplicity in the example. ShiftDiff 216 is computed based on FromFrame 210, ToFrame 212 and Offset 214. This means that the centre row in ShiftDiff corresponds to no motion, while the two

lowest rows correspond to motion one and two steps to right respectively and the two upper rows correspond to motion two and one steps to left respectively. The search range is thereby +/- 2 element positions.

Edges must be treated specially. In the given numerical example, this is taken care of by computing ShiftDiff according to the rule that when a combination of FromFrame element position, shift and Offset corresponds to a position in ToFrame outside of the given elements, then ToFrame is considered to be extended beyond its limits by repeating the outermost element.

Another solution to edges is to mark elements in ShiftDiff corresponding to combinations of element positions, shifts and offsets outside of ToFrame as undefined, and skip these values in all subsequent processing.

There may also be other reasons than edges that certain elements of ShiftDiff should be kept out of subsequent processing. One such reason could be that an estimate of motion only for a selected part of FromFrame is wanted. By marking the elements in ShiftDiff not corresponding to this selection as undefined and thereby skipping these in subsequent processing, unwanted information from other parts of FromFrame will not be propagated into the selected part. This can also give processing speed or memory consumption advantages.

Another reason for keeping certain elements of ShiftDiff out of subsequent processing is that information on occlusions may be known. An occlusion is defined as two parts of FromFrame moving so that in ToFrame, the first part will cover the second part. Assume that motion, and thereby also its position in ToFrame, is known for such a first part. When motion for the second part is to be analyzed, the shifts that would result in elements of the second part to be moved to positions that are already occupied by elements for the first part will not contribute well to the solution.

One way of achieving this is to let each element in ToFrame have with it an associated value indicating if this element is considered to be a part of an already known, occluding part. Combinations of element positions, shifts and possibly offsets leading

to elements of FromFrame ending up in the occluded part of ToFrame should then be handled by marking the corresponding element in ShiftDiff as undefined.

For further details reference is made to the article "Motion Estimation with Detection of Occlusion Areas" by Denommier/Dubois (IEEE, 0-7803-953209/92) which is incorporated herein by reference.

With prior art, one way of proceeding estimating movement would be to select the lowest value from each column in ShiftDiff 216. If the lowest value is duplicated in some column, there must be one or more guidelines for selecting which to use; one possible guideline is to select the first found, assuming that ShiftDiff is searched from top and downwards. According to these rules, shifts 218 can be selected, and the positions of the selected shifts then directly correspond to a motion field 220.

The motion field 220 in the numerical example indicates some of the problems according to this simplistic method: In some areas of the image, there will be ambiguities, since there will be several possible motions that will result in the same measurement of fit. E.g. the element with value 8 in FromFrame 210 has been matched with the first occurrence of an element with value 8 in ToFrame, though it seems based on its neighbours that it should rather have been matched with the second occurrence of the value 8 in ToFrame 212. The motion field is not smooth, as would have been expected for many types of sensors, e.g. TV cameras recording data from the interior of rigid convex objects.

The known methods of analyzing have not utilized any possible side information available from other sources. There has been no attempt to estimate motion in such a way that it will fit to a model. The present invention attempts to solve these problems, as will be shown in the following.

There may be side information available. Where this side information comes from is known to the expert, but possible sources could be previous analysis in lower resolution, or a prediction based on previous frames.

Each such set of side information is called a "hypothesis", in short Hyp. Each Hyp specifies one particular motion field. As mentioned early in the description of this first preferred embodiment, this example has a hypothesis of no motion in the right part of FromFrame. This is shown in more detail in Fig. 2 Hyp 230.

Hyp may also have associated with it information that further specifies the statistical properties of Hyp.

One such associated information to Hyp is its strength, called HypStrength. HypStrength indicates how strongly Hyp should influence the final result. Its value can be given as one value to be applied for all samples, or one value associated with each sample or each block of samples. One such HypStrength 232 is shown in Fig. 2.

Each element in Hyp 230 can be seen to select one element in ShiftDiff 216: The position of the element in Hyp selects a column in ShiftDiff, and the value of the element selects a row in ShiftDiff. This element in ShiftDiff should then be adjusted according to the HypStrength. The method of adjusting chosen for this preferred embodiment is to subtract according to HypStrength:

For each Element in Hyp:

$$\begin{aligned} \text{ShiftDiff}[\text{Hyp}[\text{Element}], \text{Element}] &= \\ \text{ShiftDiff}[\text{Hyp}[\text{Element}], \text{Element}] - \text{HypStrength}[\text{Element}] \end{aligned}$$

This is called "rewarding" a hypothesis.

Usually, Hyp 230 has been estimated with only limited accuracy, so instead of just rewarding one shift, a whole neighbourhood of shifts around the selected element of ShiftDiff 216 should be rewarded. A simple realization of this invention could use a fixed neighbourhood range. A realization that is computationally more demanding, but produces better results, can be obtained by also exploiting information on how precise Hyp 230 has been estimated in different parts of the frame. For this, a HypRange 234 can be used as a specifier of the statistical properties of Hyp. This specifies the range to the most remote neighbouring shift that should also be rewarded. The reward should

be a non-increasing function of the distance. For a function decreasing linearly with distance, this would be the definition of the rewarding mechanism:

For each Element in Hyp:

For each Neighbour inside HypRange[Element] within Dist1D from Element:

ShiftDiff[Neighbour,Element] =

ShiftDiff[Neighbour,Element] -

HypStrength[Element] *

(HypRange[Element] -

Dist1D(Neighbour,Element))/HypRange[Element]

"Dist1D" is a function that returns the difference between two elements. For 1D frames it can have this definition:

Dist1D(a,b):

abs(a-b)

There could have been one HypRange for leftwards motion and one HypRange for rightwise motion. A HypRange 234 is exemplified in Fig. 2.

HypStrength 232 can also have negative values. This corresponds to a motion hypothesis which is unwanted for some reason, possibly because it has been tested before and found to be unfruitful.

In this preferred embodiment, both Hyp 230, HypStrength 232 and HypRange 234 are given such that for each of them, there is a one to one correspondence between their elements and the elements in FromFrame 210. Another possibility would have been to have some or all of them represented in lower resolution than FromFrame. In this case, they would have to be converted into the same resolution as FromFrame before or at the stage of adding rewards to ShiftDiff. If Hyp, HypStrength or HypRange were represented with only blockwise values, this conversion would consist of duplicating values. Other methods like interpolation could be used.

If there are several sets of side information available, the rewarding procedure can be repeated for each hypothesis.

Fig. 2 shows ShiftDiff 236 after rewarding with Hyp 230 according to HypStrength 232 and HypRange 234. The selected shifts 238 and the corresponding motion field 240 are also shown. It can be seen that rewarding according to Hyp 230 has indeed affected the estimate of the motion field in that the leftmost part of the motion field now contains zeros, corresponding to no motion.

Because of the well known aperture problem, motion cannot be well estimated by just analyzing each element separately. To overcome this problem, information is propagated between the elements. In the first preferred embodiment, this is achieved by blurring ShiftDiff 236 along the element dimension.

There may be reasons for propagating information differently in different parts of the frames. E.g, in one part of the frame there may be large uniform areas, where information needs to be propagated in from the nearest textured area, while in other parts of the frame, there may be sufficient details to allow good motion estimation based only on a small neighbourhood, and information from other parts of the image would only decrease the quality of the estimate. To account for this, there is a PropRange 250 associated with each element of FromFrame. Each element of PropRange 250 indicates how information should be propagated.

The propagation can be done by convolving ShiftDiff 236 with a filter:

For each Shift in ShiftDiff:

For each Element along this Shift:

Sum = 0

SumWgt = 0

For each Neighbour within PropRange[Element] of Element:

Wgt = (PropRange[Element]-Dist1D(Neighbour,Element))
/ PropRange[Element]

Sum = Sum + Wgt * ShiftDiff[Shift, Neighbour]

$$\text{SumWgt} = \text{SumWgt} + \text{Wgt}$$
$$\text{ShiftDiffBlurred}[\text{Shift}, \text{Element}] = \text{Sum} / \text{SumWgt}$$
$$\text{ShiftDiff} = \text{ShiftDiffBlurred}$$

The above PropRange defines how far information should be propagated. Another possible way of controlling propagation is to let each element or each block of elements in FromFrame have with it an associated propagation strength, called PropStrength in short, that defines how strongly the element or block of elements should have its information propagated relative to the other elements or block of elements. The propagation taking PropStrength into consideration can be done as follows:

For each Shift in ShiftDiff:

For each Element along this Shift:

$$\text{Sum} = 0$$
$$\text{SumWgt} = 0$$

For each Neighbour within PropRange[Element] of Element:

$$\text{Wgt} = (\text{PropRange}[\text{Element}] - \text{Dist1D}(\text{Neighbour}, \text{Element}))$$
$$/ \text{PropRange}[\text{Element}] * \text{PropStrength}[\text{Neighbour}]$$
$$\text{Sum} = \text{Sum} + \text{Wgt} * \text{ShiftDiff}[\text{Shift}, \text{Neighbour}]$$
$$\text{SumWgt} = \text{SumWgt} + \text{Wgt}$$
$$\text{ShiftDiffBlurred}[\text{Shift}, \text{Element}] = \text{Sum} / \text{SumWgt}$$
$$\text{ShiftDiff} = \text{ShiftDiffBlurred}$$

In Fig. 2, a ShiftDiff 252 is shown after propagating information in ShiftDiff 236 from the reward step according to PropRange 250, together with the selected shifts 254 and the corresponding motion field 260, which is the final result.

The above method shows how to propagate information between elements in ShiftDiff after that the measurements of fit have been adjusted according to hypotheses. Another method that can be used instead of or in addition to this is to let each hypothesis have associated with it a propagation range, called HypPropRange in short. This can be one value for the whole FromFrame, or it can be one value for each

element or block of elements in FromFrame. HypPropRange can be used for propagating the effect of each hypothesis independently among the elements.

Using offsets, there are two different ways of propagating information through ShiftDiff 252: Either the propagation can be done as above, addressing ShiftDiff directly through its Shift and Element addresses. This promotes a motion field which has the same smoothness characteristics as Off. Or the propagation can be modified to take Off into account. This promotes a motion field whose smoothness is much less affected by Off. A definition of this follows:

For each Shift from min(Off)-SearchRange to max(Off)+SearchRange:

For each Element along this Shift:

Sum = 0

SumWgt = 0

For each Neighbour element within PropRange[Element] of Element:

If Shift >= Offset[Neighbour] - SearchRange and

Shift <= Offset[Neighbour] + SearchRange:

Wgt = (PropRange[Element] - Dist1D(Neighbour, Element))
/ PropRange[Element]

Sum = Sum + Wgt * ShiftDiff[Shift+Off[Element], Neighbour]

SumWgt = SumWgt + Wgt

If SumWgt > 0

ShiftDiffBlurred[Shift, Element] = Sum / SumWgt

else

ShiftDiffBlurred[Shift, Element] = Undefined

ShiftDiff = ShiftDiffBlurred

"min" is a function returning the minimum value among several values; "min(Off)" means the smallest of the Off values defined for the elements. "max" is similar, but returns the maximum value. "Undefined" is a special reserved value that can be recognized and thereby skipped in subsequent steps.

The motion field resulting from the above operations can be analysed to produce a difference field, or Diff for short. Diff is the difference between each element in FromFrame and the element in ToFrame corresponding to the selected shift:

Diff:

For each Element in FromFrame:

Diff[Element] =

abs(FromFrame[Element] - ToFrame[Element + Shift[Element]])

This is called Displaced Frame Difference, or DFD for short, in some prior art.

It is sometimes also of interest to analyze the sensitivity of the found solution: How much can the estimate of the motion be changed without producing a too bad Diff ? For this, a "Slack" is defined.

Slack can be computed in each direction separately; one part, SlackLeft 272, can indicate how much the estimate could be changed in the leftwards direction and another part SlackRight 274 can indicate the same in rightwards direction. Alternatively, a simplified Slack can be computed as the sum, average, minimum or maximum of SlackLeft and SlackRight.

For further explanation, in Fig. 6, the computation of slacks is shown. A curve 610 indicating lack of fit 620 as a function of candidate shifts 630 is shown. Based on the smallest lack of fit 640, a shift has been selected 650. A SlackLevel 660, given as input, is then used for determining SlackLeft 670 for this element as the distance between the selected shift and the first point of the lack of fit curve on the left side of the selected shift where the lack of fit exceeds SlackLevel. SlackRight 680 is found by a similar analysis on the right side of the selected shift.

In Fig. 6, the lack of fit is drawn as a continuous curve, implying that a lack of fit could be defined also at positions not corresponding to discrete shifts. This can be achieved by interpolating between evaluated shifts.

Another possibility is to only consider discrete shifts when calculating slack. This can be done according to the following:

For each Element in FromImage:

SlackLeft[Element] = 0

While $\text{abs}(\text{FromFrame}[Element] - \text{ToFrame}[Element] - \text{SlackLeft}[Element])$

< SlackLevel and $\text{SlackLeft}[Element] \leq \text{SearchRange}$:

$\text{SlackLeft}[Element] = \text{SlackLeft}[Element] + 1$

SlackRight[Element] = 0

While $\text{abs}(\text{FromFrame}[Element] - \text{ToFrame}[Element] - \text{SlackRight}[Element])$

< SlackLevel and $\text{SlackRight}[Element] \leq \text{SearchRange}$:

$\text{SlackRight}[Element] = \text{SlackRight}[Element] + 1$

Referring again to Fig. 2, based on a ShiftDiff 252, those elements 270 that are within a SlackLevel chosen to be 2 give rise to SlackLeft 272 and SlackRight 274.

Slack can be generated by analyzing values from ShiftDiff at and around the positions of the selected shift for each element. ShiftDiff at various stages can be used, each with its own advantages:

Diff or Slack can be generated based on the raw ShiftDiff 216. This has the possible advantage that differences according to the real input frames are computed.

Diff or Slack can be generated based on ShiftDiff after rewarding 236. This has the possible advantage that more information is taken into consideration, and thereby more "correct" results are computed.

Diff or Slack can be generated based on ShiftDiff after propagation 252. This has the possible advantage of better immunity to noise, since single elements are now of less importance for the computation.

Second preferred embodiment

Now the method described under first preferred embodiment will be described as applied on one type of 2D frames, namely images. The elements of the frames will now be called pixels, according to tradition in image processing.

According to the introduction of a new dimension, some of the entities defined under the first preferred embodiment will now have a different format:

Elements of FromFrame and ToFrame will now be addressed using two coordinates instead of one, the first being the vertical and the second being the horizontal position inside the respective frames.

Offsets will now for each element be given as one vertical offset and one horizontal offset respectively.

A hypothesis will for each element specify both the vertical and horizontal shift.

A hypothesis range may for each element be specified as four values, one for each of the directions up, down, left and right respectively, or alternately as two simplified values, one for vertical and one for horizontal directions, or even as one value summarizing all directions.

For images, ShiftDiff can be defined as follows:

ShiftDiff:

For each vertical position (i.e. scan line), v, in FromFrame:

For each horizontal position (i.e. pixel) on the scan line, h, in FromFrame:

For each vertical shift within SearchRange, dv:

For each horizontal shift within SearchRange, dh:

$ShiftDiff[v, h, dv, dh] = \text{abs}(\text{FromFrame}[v, h] - \text{ToFrame}[v + dv, h + dh])$

The method is described as applied to frames where each pixel has one component, like gray level television used in former times. It should be noted, however, that these gray level frames may have been preprocessed as mentioned in the first preferred

embodiment. Such preprocessing may include conversion from color, e.g. RGB, to a single component, e.g. Y as defined in PAL or NTSC color television standards.

Another preferred embodiment is to let both FromFrame and ToFrame have color components, e.g. Red, Green and Blue, or Y,I,Q or Y,U,V as defined in PAL and NTSC respectively. The entries in ShiftDiff can then be computed according to a 3-dimensional distance function, e.g. Euclidian distance, city-block distance, distance according to a uniform color space, or similar. This allows estimating motion in color images also in those regions where the intensity is constant and only the mixture of colors change.

For 2D frames, the equivalent of the "Dist1D" function used for rewarding in the 1D case can be defined as Euclidian distance:

Dist2D(x1,y1, x2,y2):

$$\sqrt{(x2-x1)^2 + (y2-y1)^2}$$

Alternatively, "Dist2D" can be defined as City-block distance:

Dist2D(x1,y1, x2,y2):

$$|x2-x1| + |y2-y1|$$

For images, especially when large search ranges are wanted, the amount of data involved becomes large. It may therefore be efficient to establish a pipelined structure. It may also become efficient to use a coarse-to-fine architecture for limiting the amount of different shifts tried. For further details reference is made to the article of W. Enkelmann "Investigations of Multigrid Algorithms for Estimation of Optical Flow Fields in Image Sequences" (Computer Vision, Graphics and Image Processing, 43, 1988) which is incorporated herein by reference. Fig. 3, 4 and 5 together show one embodiment of the combination of these, where Fig. 3 shows the overall organization working over several resolutions, Fig. 4 a module for operating on one resolution in Fig. 3, and Fig. 5 shows details of the reward steps 444, 450 from Fig. 4.

A FromFrame 310 and a ToFrame 312 each pass through a series of reduce steps 320, 322, 324, 326 where in each step the resolution is halved by a factor of two. Likewise a hypothesis Hyp 314 passes through similar steps, where the resolution of both its vertical and horizontal component is in each steps is halved by a factor of two.

The number of reduce steps in Fig. 3 is two for each frame and for the hypothesis, so that the system works on three different resolutions, but another number of steps could have been used.

EstMov 340, which corresponds to EstMov 101 in fig. 1b, estimates motion between the two frames of lowest resolution 376 and 372, also based on side information from the reduced version of Hyp 356 together with reduced versions of HypRange 358 and HypStrength 360. The result is a motion estimate, a slack, and a difference. All three are expanded in resolution by a factor of 2 in the Expand step 342. The resulting motion estimate is used both as an offset Off and a hypothesis Hyp to the next EstMov 344, which works at an intermediate resolution. The slack is used as a hypothesis range. The difference is transformed, e.g. by taking its reciprocal value, before it is fed to EstMov 344 as a hypothesis strength. This EstMov 344 also receives intermediate resolution versions of Hyp 350, HypRange 352 and HypStrength 354. The result is again expanded in resolution by an Expand 346 module, and the result is used both as an offset, hypothesis, hypothesis range and hypothesis strength by the EstMov module 348, together with full resolution versions of Hyp 314, HypRange 316 and HypStrength 318.

One of the shown EstMov modules 340 has one hypothesis as input, while the other two 344, 348 must be able to handle two input hypotheses. An EstMov capable of handling the latter case is shown in Fig. 4, and it can also handle the first case by setting HypStrength to zero for the last hypothesis.

In Fig. 4, the processing is organized around a pipeline, shown by its segments 434, 438, 442, 448, 454, 458.

A ToFrame of some resolution is stored in the ToFrame store 432, at least partially prior to other processing. Then the processing of pixels begins. A v,h looper 428 generates vertical (v) and horizontal (h) base addresses that loops over all pixels in FromFrame.

Synchronous to this, the offset Off 414 is fed to the system, and the sum of base addresses and offsets are added together in 430. The result is used for addressing into ToFrame store 432. A neighbourhood of pixels around the addressed position in ToFrame is then output from ToFrameStore 432 onto the pipeline 434.

Synchronous to the pixel neighbourhoods on the pipeline 434, pixels of FromFrame 410 arrive. From each of the pixels in the neighbourhood, the Subtractor 436 subtracts the value of the current pixel in FromFrame 410. The Absolute value module 440 then takes the absolute value of each of the neighbourhood differences. It can now be seen that this has produced a representation of ShiftDiff, where the two Shift dimensions corresponds to the neighbourhood on the pipeline and the two frame dimensions corresponds to processing time. For each time step, a new pixel of FromFrame enters the pipeline and each of the further steps of the pipeline performs one set of operations corresponding to one pixel of FromFrame.

This ShiftDiff 442 can now be rewarded according to side information. In Fig. 4, only two such reward stages 444, 450 are shown, but these structures can be repeated according to the number of different side informations available. Off 414, delayed by a delay module 446 to compensate for processing time in intermediate steps 430, 432, 436, 440, is fed to a Reward module 444. A first hypothesis Hyp1 416 with associated range HypRange1 418 and strength HypStrength1 420 are also fed to the same module. The same process is repeated for a second hypothesis Hyp2 422 in a second Reward 450 module. The result is propagated among pixels by the Propagate module 456. Shifts are selected, and Diff 472 and Slack 474 are computed, by the SelectShift module 460. The selected shifts 470 together make up the motion field.

A more detailed description of the Reward 444, 450 steps from Fig. 4 is given in Fig. 5.

The Reward module 500 is drawn as if it were designed for 1D frames, but its 2D frame equivalent can be found just by replacing all of its one component signals by two component signals, one for each dimension.

A value 514 corresponding to a center row element in ShiftDiff as defined in the first preferred embodiment arrives to the the Reward module 500 together with values corresponding to neighbouring shifts in ShiftDiff 510, 512, 516, 518.

In the Figure, five groups of processing elements can be seen, each corresponding to one shift. One of these groups will now be explained, with the understanding that the other perform similar functions.

The corresponding image coordinate (or coordinates, for the 2-D case) of these values is now found by adding the common offset value Off 520 to a number (or two numbers for the 2-D case) 530 indicating the position of the neighbouring shifts relative to the center value 514.

The distance between the image coordinate and the hypothesis Hyp 522 for each pixel is then found by computing absolute value of the difference between them (or by applying one of the Dist2D functions defined earlier for the 2-D case).

The fraction of HypStrength that should be used for adjusting the measurement of fit is computed by a subtract module 538 and a divide module 540. The subtract module subtracts the distance between the image coordinate and the hypothesis. This subtract is a so called zero clamped operation, in that any resulting value below zero will be set ("clamped") to zero. The divide module devide the result of the subtraction with HypRange 524.

The result is multiplied by HypStrength 526 before beeing subtracted from the input 510, thereby producing an adjusted value 550 corresponding to an element of an updated shift difference matrix.

It can be seen that this cascade of processing elements in total implements the following expression:

$$\begin{aligned} \text{ShiftDiff[Shift,Element]} &= \text{ShiftDiff[Shift, Element]} - \\ &\text{HypStrength[Element]} * \\ &(\text{HypRange[Element]} - \text{abs}(\text{Off[Element]} + \text{Shift} - \text{Hyp[Element]})) \\ &/ \text{HypRange[Element]} \end{aligned}$$

Each of the processing elements in Fig. 5 can operate with low precision and no memory is needed, which gives freedom in choice of implementation. Each processing element could be implemented as a digital arithmetic operation, using a digital look-up table covering only a few selected possible input values, or using analog electronics either constructed by Operational Amplifiers or directly by transistors.

Third preferred embodiment

In the first preferred embodiment, the invention was shown as applied to 1D frames. There ShiftDiff had one dimension corresponding to the dimension inherent in the frame structure and one dimension according to the dimension for which the motion was to be analyzed, namely the same dimension.

In the second preferred embodiment, it was shown how to apply the invention to 2D frames. ShiftDiff had two dimensions corresponding to the dimensions of the frame structure and two more dimensions according to the dimensions for which the motion was to be analyzed. It was shown how the hypotheses had ranges in the various dimensions, how offsets could be used in certain dimensions, and how information could be propagated in various dimensions.

Likewise, this can be generalized to 3D frames. For an example, consider whole-body X-ray computed tomography scans of a human. In such systems, the scan is often performed in time steps such that one "slice" of the body is scanned, then the scanning mechanisms is advanced one step along some axis, and this is repeated until the whole body has been analyzed. The analysis of such a scan for some purposes

depend on alignment of the slices. This can be difficult due to breathing or other motions by the subject. By estimating and compensating the motion between slices, the effect of the subject motion can be made less severe.

The motion may be assumed or known to occur in less than all dimensions inherent in the frame structure. E.g. for a human laying on a bench, chest motion caused by breathing will mainly occur in the directions perpendicular to the main axis of the body, which is the same axis as that used for stepping. In this case, ShiftDiff will have two dimensions corresponding to a slice, one dimension corresponding to the number of slices, and two dimensions according to the dimensions where the motion is assumed to happen, that is, within each slice.

Side information about the motion of some parts of the subject could be obtained by ultrasound sensors, analysis of traditional TV camera images synchronized with the scanning, or by other traditional sensors. These would then provide hypotheses that could be used to improve the estimation of motion. The hypothesis would be highly valid and thereby have high strengths for parts easily accessible to the side information sensors, e.g. the body surface or localized markers, and less valid for other parts. Information could then be propagated through the wanted dimensions.

Likewise, 3D MRI scans from the brain of a set of persons can be more easily and precisely interpreted when position and shape information has been separated from chemical compositional information. The position and shape normalization can be done by the 3D motion estimation and corresponding compensation.

While the invention has been particularly shown and described with reference to preferred embodiments thereof, it will be understood by those skilled in the art that various changes in form and detail may be made therein without departing from the spirit and scope of the invention.

CLAIMS

1. A method for estimating motion between at least two frames, each consisting of a plurality of samples of an input signal, the method comprising the steps of:

- (1) selecting for each sample in the first frame a plurality of candidate shifts,
- (2) computing for each candidate shift a measurement of fit between said sample in said first frame and the corresponding sample at the shifted position in said second frame, and
- (3) selecting for each sample in said first frame the shift with the best associated fit,

wherein the collection of selected shifts represents the motion from said first frame to said second frame,

characterized in that,

prior to step (3),

- (4) for each sample in said first frame the measurement of fit is adjusted according to an externally available hypothesis indicating a priori knowledge about how probable it is that certain shifts represent the correct motion.

2. The method according to claim 1 wherein the motion between the first frame and two other frames (second and third frames) is estimated by

selecting for each sample in the first frame a plurality of candidate shifts in the two other frames,

computing for each candidate shift the measurement of fit both between said sample in said first frame and the corresponding sample at the shifted position in said second frame and

said sample in said first frame and the corresponding sample at a shifted position in said third frame according to a scaled version of said shift between said first and second frames, said scaling being indicative of a priori knowledge about similarity of said frames.

3. The method according to claim 1 or 2 wherein the measurement of fit is adjusted according to more than one externally available hypothesis.

4. The method according to one of claims 1 to 3 wherein each sample or block of samples in each hypothesis has associated with it a measurement of how probable it is that certain shifts represent the correct motion for said sample or block of samples, and the amount each measurement of fit is adjusted is a function of these probabilities.

5. The method according to one of claims 1 to 3 wherein each sample or block of samples in each hypothesis has associated with it a measurement of how precise the hypothesis is with regard to which shifts it applies for, and that for each sample not only the measurement of fit for the sample directly indicated by the hypothesis, but also the measurement of fit similar shifts are adjusted.

6. A method for estimating motion between at least two frames each consisting of a plurality of samples of an input signal, the method comprising the steps of:

(1) selecting for each sample in the first frame a plurality of candidate shifts,
(2) computing for each candidate shift a measurement of fit between said sample in said first frame and the corresponding sample at the shifted position in said second frame, and

(3) selecting for each sample in said first frame the shift with the best associated fit,

wherein the collection of selected shifts represents the motion from said first frame to said second frame,

characterized in that,

prior to step (3),

(4) for each sample in said first frame the measurement of fit for each shift is adjusted according to measurements of fit for the same shifts for adjacent samples in said first frame.

7. The method according to claim 6, wherein the process of adjusting measurements of fit for shifts according to measurements for the same shifts for adjacent samples is carried out by convolving the same shifts for adjacent samples with a low pass filter.

8. The method according to one of claims 6 or 7, wherein for each hypothesis, a value associated with each sample or block of samples controls how far away adjacent samples will have their measurements of fit for shifts adjusted.

9. The method according to one of claims 6 or 7, wherein a value associated with each sample or block of samples in one of the frames controls how far away adjacent samples will have their measurements of fit for shifts adjusted.

10. The method according to one of claims 6 to 9, wherein a value associated with each sample or block of samples in one of the frames controls how strongly adjacent samples will be influenced.

11. The method according to one of claims 1 to 10, wherein with each sample or block of samples in said first frame there is an associated value that controls the selection of candidate shifts.

12. The method according to one of claims 1 to 11, wherein steps (1) to (3) are executed once for each of several resolutions of the frames, starting from the lowest resolution, and wherein, for each resolution other than the lowest, the measurement of fit is adjusted according to the found estimate of motion for the lower resolution.

13. The method according to claim 1, wherein after the steps (1) to (3), the found shifts are analyzed with regard to uniqueness, and the result of the analysis is used to control the adjustment of measurements of fit for a subsequent application of steps (1) to (3).

14. The method according to claim 1, wherein
after the steps (1) to (3), the found measurements of fit are used to control the
adjustment of measurements of fit for a subsequent application of steps (1) to (3).

15. An apparatus for estimating motion between at least two frames each
consisting of a plurality of samples of an input signal, the apparatus comprising:
means for selecting for each sample in the first frame a plurality of candidate
shifts,

means for computing (128) for each candidate shift a measurement of fit
between said sample in said first frame and the corresponding sample at the
shifted position in said second frame, and

means for selecting (144) for each sample in said first frame the shift with the
best associated fit,

wherein the collection of selected shifts represents the motion from said first
frame to said second frame,

characterized by

means, coupled between said means for computing (128) and said means for
selecting (144), for adjusting (132, 136) for each sample in said first frame the
measurement of fit according to an externally available hypothesis indicating a
priori knowledge about how probable it is that certain shifts represent the correct
motion.

16. An apparatus for estimating motion between at least two frames each
consisting of a plurality of samples of an input signal, the apparatus comprising:

means for selecting for each sample in the first frame a plurality of candidate
shifts,

means for computing (128) for each candidate shift a measurement of fit
between said sample in said first frame and the corresponding sample at the
shifted position in said second frame, and

means for selecting (144) for each sample in said first frame the shift with the
best associated fit,

wherein the collection of selected shifts represents the motion from said first frame to said second frame,

characterized by

means, coupled between said means for computing (128) and said means for selecting (144), for adjusting (140) the measurement of fit for each shift according to measurements of fit for the same shift for adjacent samples in said first frame.

17. The apparatus according to claim 15 or 16, wherein, coupled between said means for computing (128) and said means for selecting (144)

means for adjusting (132, 136) for each sample in said first frame the measurement of fit according to an externally available hypothesis indicating a priori knowledge about how probable it is that certain shifts represent the correct motion and

means for adjusting (140) the measurement of fit for each shift according to measurements of fit for the same shift for adjacent samples in said first frame are provided for.

18. The apparatus according to claims 15 to 17, adapted to be used according to the method of anyone of claims 1 to 14.

1/6

Fig. 1a

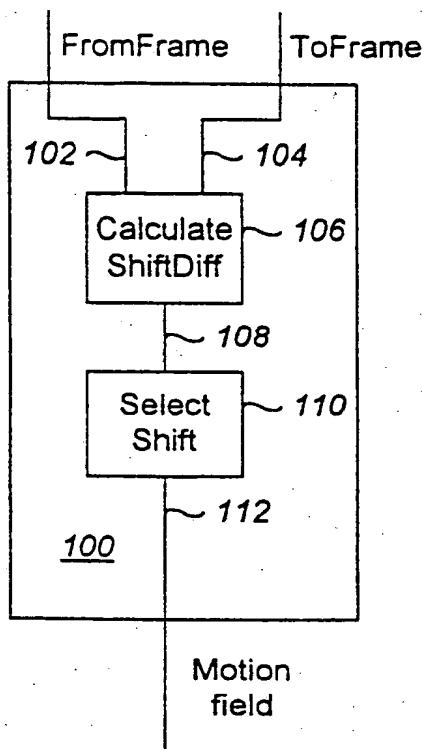


Fig. 1b

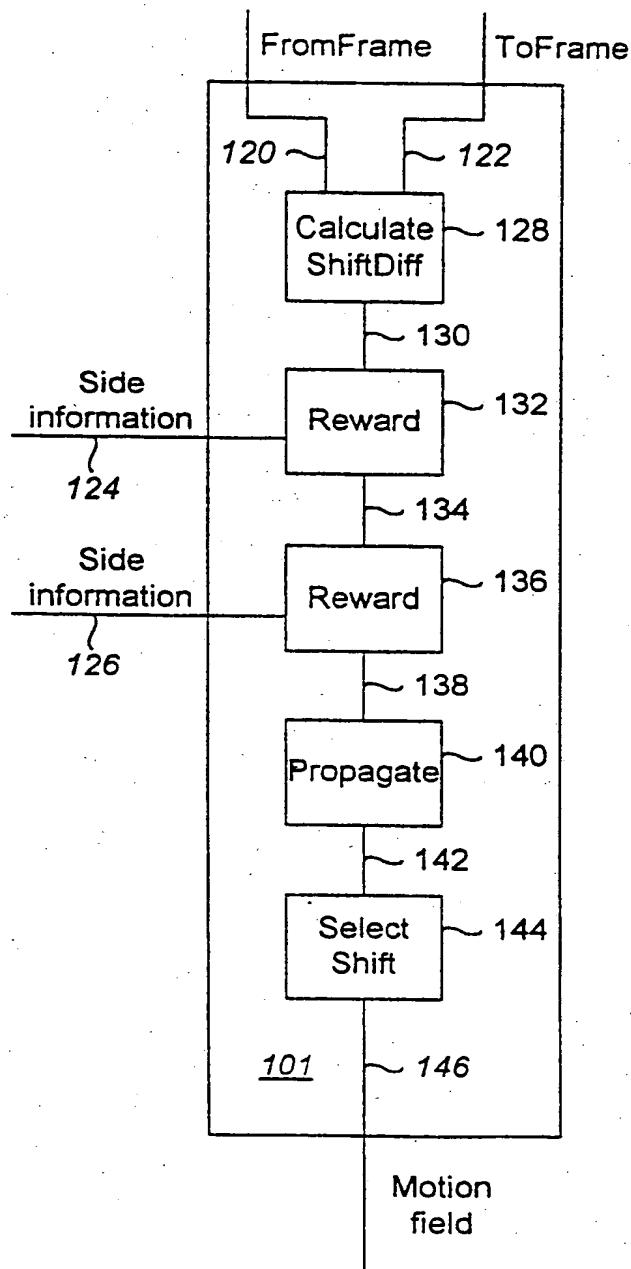
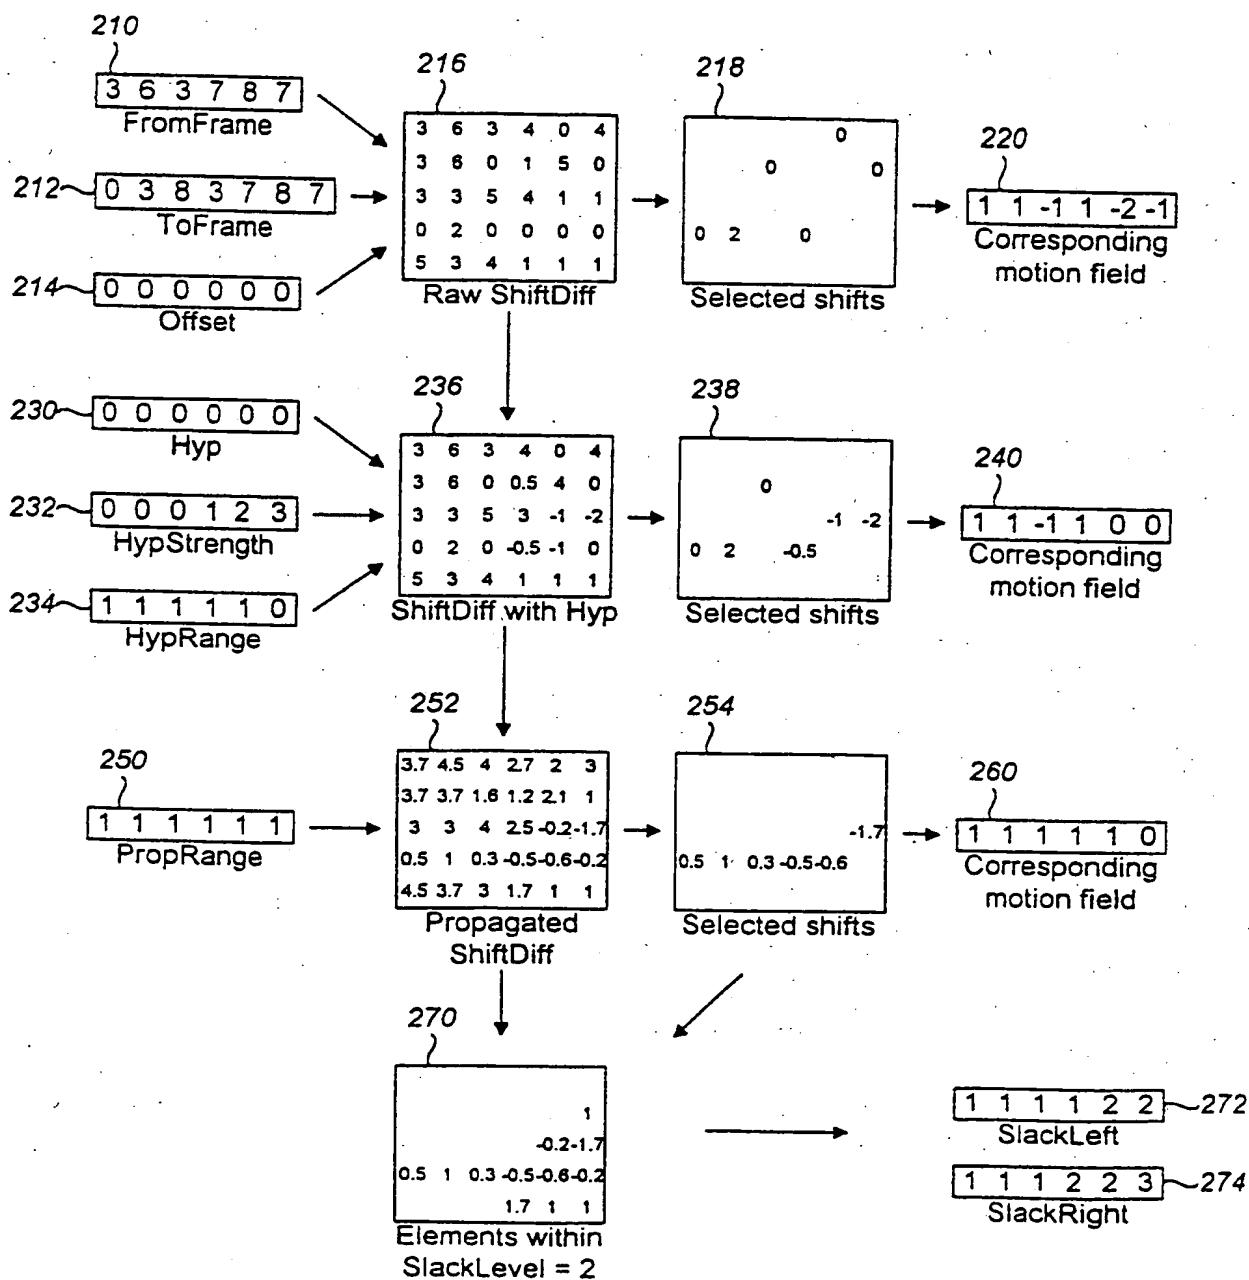
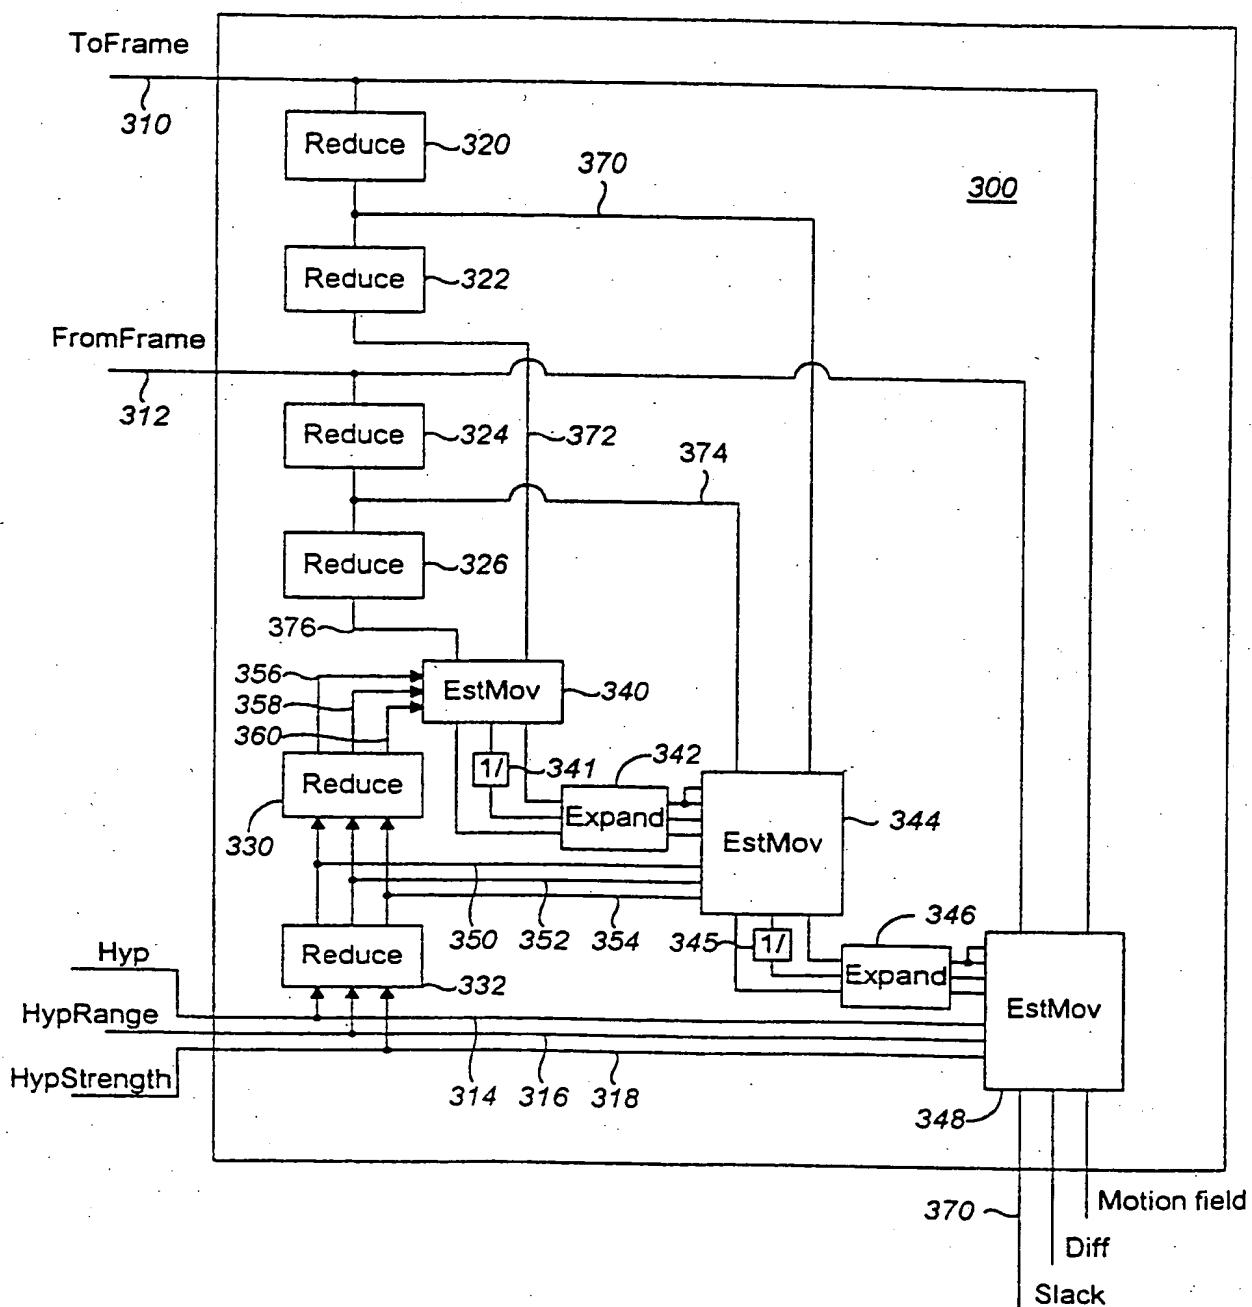


Fig. 2



3/6

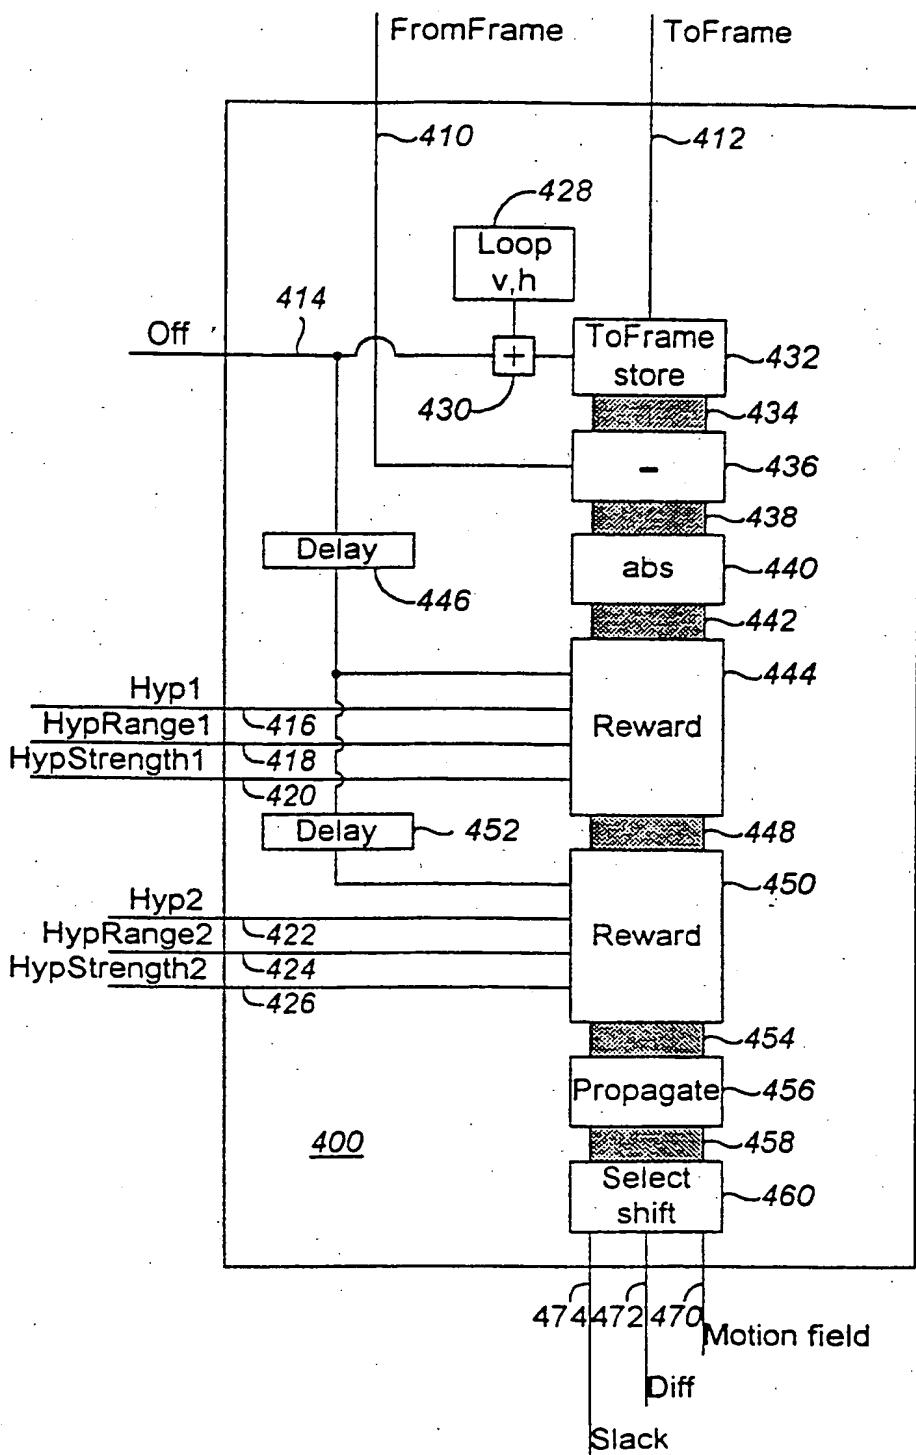
Fig. 3



SUBSTITUTE SHEET (RULE 26)

4/6

Fig. 4



SUBSTITUTE SHEET (RULE 26)

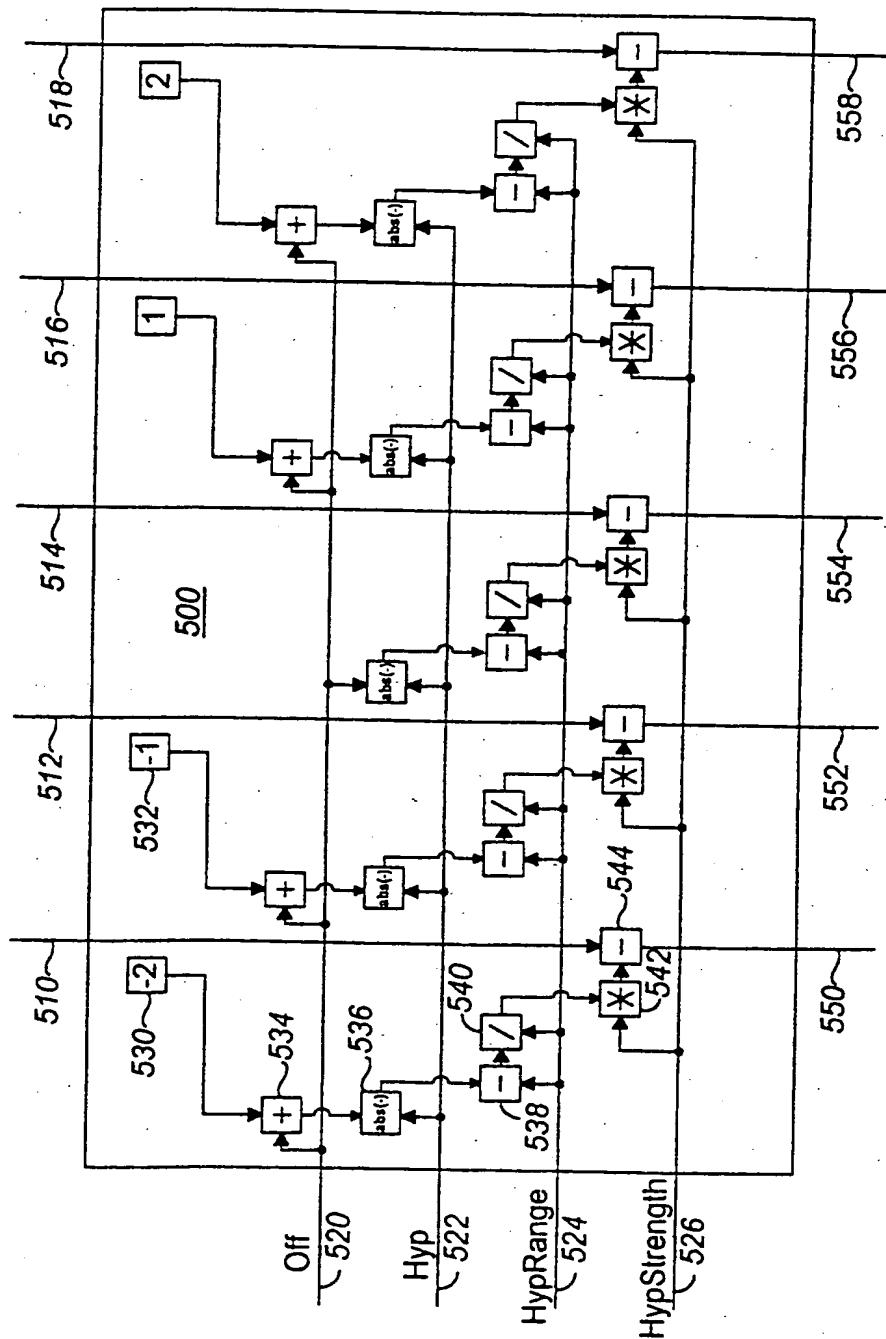
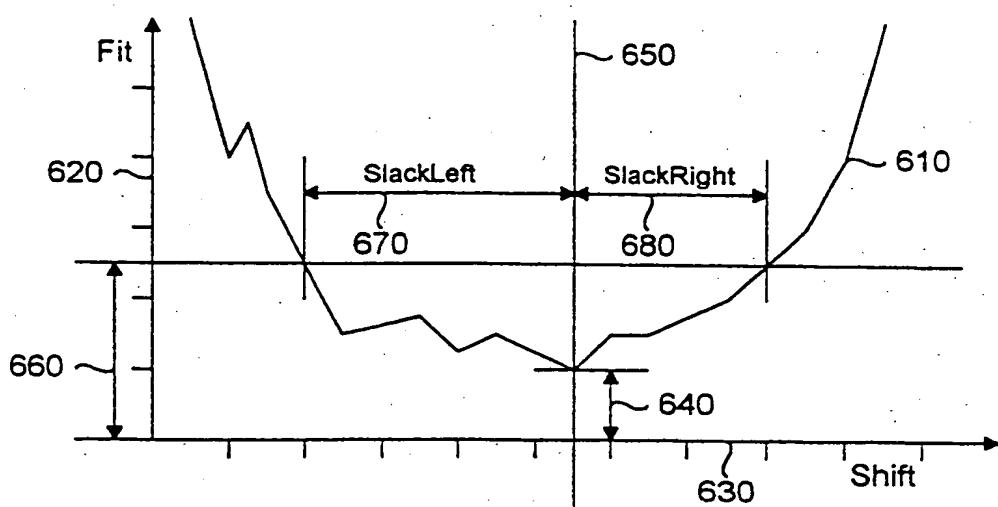


Fig. 5

6/6

Fig. 6



INTERNATIONAL SEARCH REPORT

Application No
EP 95/01011A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 G06T/20

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 6 G06T

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP,A,0 498 938 (TELENORMA GMBH) 19 August 1992 see abstract ---	6,16
A	EP,A,0 362 849 (KABUSHIKI KAISHA TOSHIBA) 11 April 1990 see abstract ---	1,6
A	EP,A,0 392 671 (MATSUSHITA ELECTRIC INDUSTRIAL CO., LTD) 17 October 1990 see abstract -----	1,5,6,16

 Further documents are listed in the continuation of box C. Patent family members are listed in annex.

* Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *&* document member of the same patent family

1

Date of the actual completion of the international search:

13 June 1995

Date of mailing of the international search report

22.06.95

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patendaan 2
NL - 2280 HV Rijswijk
Tel. (+ 31-70) 340-2040, Tx. 31 651-epo nl,
Fax: (+ 31-70) 340-3016

Authorized officer

Chateau, J-P

INTERNATIONAL SEARCH REPORT

Information on patent family members

J Application No

[REDACTED] /EP 95/01011

Patent document cited in search report	Publication date	Patent family member(s)		Publication date
EP-A-498938	19-08-92	DE-A-	4104067	13-08-92
EP-A-362849	11-04-90	JP-C-	1789131	29-09-93
		JP-A-	2100583	12-04-90
		JP-B-	4079190	15-12-92
		US-A-	5034988	23-07-91
EP-A-392671	17-10-90	JP-A-	2246686	02-10-90
		US-A-	5157732	20-10-92

THIS PAGE BLANK (USPTO)